

AMENDMENTS TO THE SPECIFICATION

Page 1, para. 1

The present invention relates to Internet information services. In particular, the present invention relates to improvements ~~related to and / or use of the Usenet. The present invention~~ also has application to email systems, as well as other electronic distribution media.

Page 1, para. 2

~~In one aspect, the present invention relates to the distribution, access and / or download speed and efficiency of relatively large binary objects, and involves a new system design and method of use.~~

Page 1, para. 3

~~In a second aspect of the present invention relates to a method that enables relatively transparent encoding within objects' URLs information necessary to locate the object in a Usenet server and retrieve it. The method also allows transparent retrieving of news cached objects from their original Web servers services.~~

Page 3, para. 2

The Usenet was originally designed for exchange of textual information, but presently the major part of bandwidth and storage resources is consumed by so called --"binary" newsgroups that mainly carry binary data. In terms of bytes, the top four newsgroups consume 22% of the entire volume. The top 35 groups consume 50% of the entire volume.

Page 3, para. 6

News articles are stored in news servers to enable users to access them. But this storage brings about another problem, that being the limited availability of storage space. To limit the amount of disk space occupied by binary newsgroups, ISPs normally set shorter expiration time ~~limit~~limits for binary postings. This helps to save disk space in short term, but users of popular binary news groups compensate for this by re-posting popular binary objects regularly, to ensure

their availability. This reduces the effect of the measures taken by ISPs and even makes the situation worse because

Page 4, para. 1

- _____1) Often a binary object is re-posted by more than one poster and this results in there being several copies of the binary object stored on the server attached to different messages, and
_____2) Regular re-posting of large binary objects is considered to lead to a waste of bandwidth that should be avoided.

Page 4, para. 5

As described above, some ISPs try to reduce expenses caused by handling binary attachments by setting a low limit on time that a message with a binary object will spend in the news pool on their server. However, this is not considered an effective solution because often the same binary object returns re-posted with a new message. This increases news feed traffic and leads to multiple copies of the same object being stored.

Page 5, para. 2

Patent No US 5,771,355 - Title: Transmitting Electronic Mail by Either Reference or Value at File-Replication Points to Minimise Costs. This patent covers technology aimed at improving e-mail delivery in certain conditions. E-mail attachments are delivered by "optimal path". For example, when the path includes intermediary points that make it much longer than the distance from the sender to the receiver, it makes sense to defer sending of the attachment until the receiver requests it and, in this case, send the attachment directly from the site where it is stored to the receiver.

Page 6, para. 4

Patent No US ~~5,815,663~~5,815,663384,565 - Title: Method and Apparatus for Identifying Duplicate Data Messages in a Communication System. This patent disclosure is considered directed at how to determine whether one message is a copy of another message in an environment where errors are very frequent. In the Usenet, however, the environment is relatively error free, and thus the problems addressed in this disclosure are not considered relevant to the problems of the present invention.

Page 7, para 1

Given the large average size of binary objects, pollution of binary newsgroups by spam and slow speed of downloading via modem lines, it is very important to help users to make better decisions on whether to download a particular binary object. Because, if this decision is wrong, they spend resources (their own time, on-line time, traffic) on downloading an object that they will discard right after downloading and examining.

Page 7, para. 3

Currently, almost the only description of an article is its subject. This way of describing information items is more or less adequate for textual messages that contain text discussing the subject. For multimedia items, one-line descriptions can hardly be adequate. Normally, the subject contains the name of the collection or short description of the multimedia item, name of file, number of the part and total number of parts (such as "Persian kitten cats123.jpg (1/1) 35567 bytes"). This format is often used, but many multimedia postings do not have even that. Often subject lines are quite meaningless, e.g. "My loved kittens".

Page 7, para. 4

A still further problem is the relatively large amount of traffic and relatively slow response times over the Internet. Users feel frustrated if they have to wait a long time for a response from their Web browser. A relatively fast response has become absolutely critical for the emerging multibillion e-commerce business. Research shows that a substantial part of users, if idle for more than 8 seconds, would exit a site without completing the transaction. Estimated \$4.8 billion is lost annually due to such bail-out behaviour.

Page 8, para. 2;

Analysis of Internet traffic shows that transmission of objects bigger than 1Mb in size takes about 40% of the total Internet traffic, which is a significant amount, considering that less than 1% of transmitted objects is this size. According to the same source, transfer error rate increases exponentially as the object size becomes larger than 10Mb and the error rate of objects larger than 10Mb is over 80%. This data shows that, firstly, large objects constitute a significant amount of Internet traffic. Thus, we can conservatively estimate that objects larger than 100K in size take at least 70% (or more) of the traffic. Secondly, this data shows that large objects are

very hard to download, not only because it is slow, but also because the process of downloading a large object is more likely to fail. This is thus considered an obstacle to the use of large multimedia objects on the Web, for example, for e-commerce and remote education services.

Page 8, beginning at para.5; through page 9

~~A first aspect of the present invention provides a method, system and / or network for transporting of Web objects from the server side (their original server) to the client side via the Usenet or a Usenet like system. The method includes:~~

~~Construction/determining/allocating a URL (Uniform Resource Locator) for the object, placing the object on the original server in such a way that this URL~~

- ~~a) — contains information necessary to find the object in a Usenet server;~~
- ~~b) — indicates that the object has been posted to the Usenet and may be found on a Usenet server; and~~
- ~~c) — that the URL can be used to retrieve the object transparently from its original server.~~

~~Furthermore, the method may include:~~

~~posting the object on the Usenet;~~

~~on the client side, intercepting requests for the object, interpreting them and using the extracted information to find the object from a Usenet server and return it to the client.~~

~~1. — A method of associating an URL with a Web object(s) for transport from a server side (their original server) to a client side via the Usenet or a Usenet like system, the method including the steps of:~~

~~—— a. — Constructing/determining/allocating a URL (Uniform Resource Locator) for the object, and~~

~~b. — placing the object on the original server in such a way that this URL~~

~~1. — contains information necessary to find the object in a Usenet server;~~

~~2. — indicates that the object has been posted to the Usenet and may be found on a Usenet server; and~~

~~—— 3. — can be used to transparently retrieve the object from its original server.~~

~~_____ This aspect also provide a method of transporting Web object(s) via Usenet, the method including:~~

~~_____ associating a URL with the Web object as outlined above,~~

~~posting the object on the Usenet;~~

~~at a client side, intercepting requests for the object, interpreting them and using information extracted, as a result of the interpretation, to retrieve the object from a Usenet server.~~

A first aspect of the present invention provides a method associating a URL (Uniform Resource Locator) with a Web object(s) for transport from a server side to a client side via a Usenet system, the method including the steps of:

on the server side,

(a) constructing and allocating said URL for the object, and

(b) placing the object on an original server in such a way that said URL

(i) contains information necessary to find the object in a Usenet server;

(ii) indicates that the object has been posted to said Usenet system and may be found on a Usenet server; and

(iii) can be used to transparently retrieve the object from said original server;

(c) associating said URL with the Web object;

(d) posting the object on the Usenet system;

and on the Client side

(e) intercepting requests for the object, interpreting them and using information extracted, as a result of the interpretation, to locate the object from said Usenet server;

(f) if the object is not found posted on a Usenet server, or its version is not current:

(i) retrieving the object from the original server;

(ii) receiving digitally signed permission to post the object on behalf of the original server and to cancel an expired version, if any, and transmitting this permission to one or more of Usenet servers along with the object.

Page 10, para. 1

~~This aspect also provides a useful method of constructing an URL useful in accordance with the method as disclosed above.~~

Page 10, para. 2

Still A further aspect of the present aspect invention provides a communication system adapted to distribute Web objects from a web host server to a client, the system having:

a Web host ~~server~~ server on which the web objects are stored, the web host server being coupled to the WWW (~~World Wide Web~~),₁

the coupling between the client, the WWW and web host server enabling bi-directional communication,

The improvement including:

providing a first Caching agent intermediate and coupled to the client and WWW and Usenet, and

providing a second Caching agent intermediate and coupled to the WWW and the Usenet and the web host server,

wherein the first Caching agent and second Caching agents enable communication of objects between the client and the Web host server to be via either the Internet or the Usenet.

wherein the first caching agent performs the following functions:

analyses Web requests containing URLs of required objects,

based on the URL, decides, whether an object has been posted to the Usenet by its original server and thus, may be found in the Usenet,

if the object has not been posted to the Usenet, the first agent passes the request further for normal processing by the Web server or cache engine,

if the object has been posted to the Usenet;

based on its configuration information, the first agent selects one or more available Usenet servers and tries to find the required object on them,

if the object is found, the first agent retrieves it and returns to the client, and / or

if the object is not available, the first agent passes the request for further processing by the original server or a caching engine;

wherein the second caching agent performs the following functions:

intercepts requests to the server and identifies those that are requesting Usenet posted objects,

if such a request is found, the second agent cleans up its URL, removing its part that concerns newsgroups or including the required information in the URL and combining it with object placement in such a way, that no further cleaning is necessary,

once cleaned, the URL is passed further for processing by the server or server side cache engine,

tracing events of modification of the server objects that are to be, or have been posted to the Usenet,

if an object has been modified (or created), the second agent cancels its previous version, if necessary in the Usenet and posts a new digitally signed one, and/or

periodically re-post objects to the Usenet to ensure their availability.

Page 10, para. 5

Thus, Usenet can be used for automatic replication and mirroring of Web objects. In the context of this task, newsgroups can be seen as subscription channels to which servers subscribe if their users are likely to retrieve posted Web objects. One of the examples could be a “Shareware channel” that would be automatically mirroring contents of Web shareware servers on the Web.

Page 11, 3rd full paragraph

A second aspect of the present invention provides a method ~~of~~ and system for creating a URL for use in the Web, the method including the steps of:

providing a first field having information sufficient to locate an object on a web server,
and

providing a second field having information sufficient to locate the object on the Usenet.

Page 11, 4th full paragraph

In essence, this aspect discloses a method and system that enables transparent encoding within objects' URLs information necessary to locate the object in a Usenet server and retrieve it. A number of example implementations are disclosed and any of these (as well as other methods as would be apparent to the skilled person) may be used in our system. These methods allow transparent retrieving of news cached objects from their original servers, in the case if that the objects could not be found in the Usenet or no Usenet server is available to the client.

Page 12, 2nd full paragraph

~~Figure 2~~ Figures 2a and 2b illustrates schematically a 1st method applicable to the present system that can be used to identify binary attachments.

Page 36, para. 8

For example, let collection number be 123, object numbers 45, 46, 47 and 48, and original subject of the message, "Cute kittens number one, two, ~~free~~ three and four". In the process of posting, the client will modify the subject in the following way, "Cute kittens number one, two, three and four id=123:45-48".

Page 39, full paragraph 5

The user ~~than~~ then can make a better informed downloading decision if they have better described articles to select from.

Page 40, para. 3

CSCA performs the following functions:

——Analyses Web requests containing URLs of required objects.

Page 40, para. 4

——Based on the URL, decides, whether an object has been posted to the Usenet by its original server and thus, may be found in the Usenet.

Page 40, para. 5

——If the object has not been posted to the Usenet, CSCA passes the request further for normal processing by the original Web server or cache engine.

Page 40, para. 6

If the object has been posted to the Usenet:

——Based on its configuration information, CSCA selects one or more available Usenet servers and tries to find the required object on them.

Page 40, para. 7

——If the object is found, CSCA retrieves it and returns to the client.

Page 40, para. 8

——If the object is not available, CSCA passes the request for further processing by the original server or a caching engine.

Page 41, para 2

3. SSCA may also periodically re-post objects to the Usenet to ensure their availability.

——The only mandatory function of SSCA is ensuring availability of the objects in the Usenet. However, this function can be performed by CSCAs on behalf of the original server, as discussed below. Thus, SSCA is not an essential element of the system, but its availability makes easier implementation of certain features: validation of objects, access control and traffic billing, without modifying Web servers.

Page 41, para. 4

——Obviously, CSCA and SSCA can be independent applications, or CSCA can be built into client and/or client side cache engine, and SSCA can be built into Web server and/or server side cache engine.

Page 41, para. 6

When a client requests an object, it must receive its current, valid version. This is not hard to ensure using validation requests in step 6 of CSCA actions. If the object is found, CSCA sends its version information, such as UBOI, to the SSCA, or a standard HTTP validation request to the original server. If the object is current, and only if, it will be ~~send~~ sent to the client. So, the problem of validation is not a hard one. Given that most Usenet cached objects are large, expenses on their validation are negligible compared to the transmission cost.

Page 43, para. 2

We will disclose below three methods that allow ~~to transparently encode~~ transparent encoding in objects' URLs information necessary to locate the object in a Usenet server and retrieve it. Any of these methods may be used in our system. These methods allow transparent retrieving of news cached objects from their original servers, in case if no CSCA is installed on the client side and no SSCA is installed on the server side. If we can assume that at least one of these agents is always installed, it can perform URL translation (for example, clear URLs of Usenet related parameters), and the problem discussed here becomes trivial.

Page 44, para. 5

Modified URL that contains information that the object has been posted to the Usenet (this conclusion ~~can~~ is made based on presence in the URL of the special string "usenetcached"), and name of its message-id is available after decoding – a process that is reverse to the process of encoding described in step 1.

Page 46, para. 1

This method ~~id~~ is based on passing the information in the query part of the URL. Because we are retrieving a file, this part would normally be ignored by Web servers. Even if we were retrieving a dynamic object that required passing query parameters, extra parameters are also normally ignored by CGI scripts processing queries.

Page 47, para. 4

When posting messages, original servers can explicitly mark them as write-protected and/or read-protected. By default, all messages are write-protected, but not read-protected. Access protection information is contained in a header invented by us ~~header~~ with a name X-Access-Protection. If this header has string “write=no”, it changes write protection of the message from the default mode. If this header has string “read=yes”, it changes read protection of the message from the default mode.

Page 48, para. 3

If the message is write-protected and the client is not the original poster or a trusted Usenet server, the server responds with code that requests a digitally signed certificate by the original poster for permission to modify the message

Page 50, para. 9

Fourth, ~~SSUCA~~-SSCA constructs a Usenet message with constructed message-id and containing a copy of the object as a binary attachment. If reading access to the object is limited, SSCA places header “X-Access-Protection: read=yes” in the message.